

**METHOD AND SYSTEM FOR NETWORK MANAGEMENT CAPABLE OF
RESTRICTING CONSUMPTION OF RESOURCES ALONG
ENDPOINT-TO-ENDPOINT ROUTES THROUGHOUT A NETWORK**

5

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to an improved data processing system and, in particular, to a method and system for multiple computer or process coordinating.

Still more particularly, the present invention provides a method and system for network management.

2. Description of Related Art

Technology expenditures have become a significant portion of operating costs for most enterprises, and businesses are constantly seeking ways to reduce information technology (IT) costs. This has given rise to an increasing number of outsourcing service providers, each promising, often contractually, to deliver reliable service while offloading the costly burdens of staffing, procuring, and maintaining an IT organization. While most service providers started as network pipe providers, they are moving into server outsourcing, application hosting, and desktop management. For those enterprises that do not outsource, they are demanding more accountability from their IT organizations as well as demanding that IT is integrated into their business goals. In both cases, "service level agreements" have been employed to contractually guarantee service delivery between an IT organization and its customers. As a result, IT teams now require management solutions that

focus on and support "business processes" and "service delivery" rather than just disk space monitoring and network pings.

IT solutions now require end-to-end management that 5 includes network connectivity, server maintenance, and application management in order to succeed. The focus of IT organizations has turned to ensuring overall service delivery and not just the "towers" of network, server, desktop, and application. Management systems must 10 fulfill two broad goals: a flexible approach that allows rapid deployment and configuration of new services for the customer; and an ability to support rapid delivery of the management tools themselves. A successful management solution fits into a heterogeneous environment, provides 15 openness with which it can knit together management tools and other types of applications, and a consistent approach to managing all of the IT assets.

With all of these requirements, a successful management approach will also require attention to the 20 needs of the staff within the IT organization to accomplish these goals: the ability of an IT team to deploy an appropriate set of management tasks to match the delegated responsibilities of the IT staff; the ability of an IT team to navigate the relationships and 25 effects of all of their technology assets, including networks, middleware, and applications; the ability of an IT team to define their roles and responsibilities consistently and securely across the various management tasks; the ability of an IT team to define groups of 30 customers and their services consistently across the various management tasks; and the ability of an IT team to address, partition, and reach consistently the managed

devices.

Many service providers have stated the need to be able to scale their capabilities to manage millions of devices. When one considers the number of customers in a 5 home consumer network as well as pervasive devices, such as smart mobile phones, these numbers are quickly realized. Significant bottlenecks appear when typical IT solutions attempt to support more than several thousand devices.

10 Given such network spaces, a management system must be very resistant to failure so that service attributes, such as response time, uptime, and throughput, are delivered in accordance with guarantees in a service level agreement. In addition, a service provider may 15 attempt to support as many customers as possible within a single network management system. The service provider's profit margins may materialize from the ability to bill the usage of a common network management system to multiple customers.

20 On the other hand, the service provider must be able to support contractual agreements on an individual basis. Service attributes, such as response time, uptime, and throughput, must be determinable for each customer. In order to do so, a network management system must provide 25 a suite of network management tools that is able to perform device monitoring and discovery for each customer's network while integrating these abilities across a shared network backbone to gather the network management information into the service provider's 30 distributed data processing system. By providing network management for each customer within an integrated system, a robust management system can enable a service provider

to enter into quality-of-service (QOS) agreements with customers.

Hence, there is a direct relationship between the ability of a management system to provide network monitoring and discovery functionality and the ability of a service provider using the management system to serve multiple customers using a single management system.

Preferably, the management system can replicate services, detect faults within a service, restart services, and

reassign work to a replicated service. By implementing a common set of interfaces across all of their services, each service developer gains the benefits of system robustness. A well-designed, component-oriented, highly distributed system can easily accept a variety of services on a common infrastructure with built-in fault-tolerance and levels of service.

Distributed data processing systems with thousands of nodes are known in the prior art. The nodes can be geographically dispersed, and the overall computing

environment can be managed in a distributed manner. The managed environment can be logically separated into a series of loosely connected managed regions, each with its management server for managing local resources. The management servers coordinate activities across the enterprise and permit remote site management and operation. Local resources within one region can be exported for the use of other regions.

A service provider's management system should have an infrastructure that can accurately measure and report

the level of service available from any given resource throughout the system, which can be quite difficult to accomplish in a large, highly distributed computing

environment. When attempting to measure the availability of a target resource, a network management task that is monitoring the target resource might use a combination of resources for the measurement operation, and the accuracy 5 of the reading may be hampered by other activities at the target resource and at other intermediate resources.

For example, a performance measurement operation on a target resource may be routed through a particular node in the network, and the node may be used by multiple 10 activities, including application-specific activities and other management activities. Any of these activities may consume a significant amount of the node's bandwidth, thereby affecting measured response times. Hence, those other activities should be restricted or terminated 15 during a performance measurement operation.

In order to fulfill quality-of-service guarantees within a network management system consisting of a million devices or more, performance measurements may be required along various network routes throughout the 20 system. Computational resources throughout the system should be controllable so that the management system can obtain accurate performance measurements along particular routes that are independent of other activities that are occurring along the route.

25 Moreover, if a service provider were able to restrict the use of resources along routes through the network for performance measurement purposes, then the service provider could restrict routes for other purposes. For example, the service provider could 30 contract with customers to provide a high level of service that requires exclusive use of a particular route.

Therefore, it would be advantageous to provide a method and system that restricts consumption of resources along particular network routes so that the performance of various target resources can be accurately measured.

- 5 It would be particularly advantageous if access to particular routes could be restricted so that particular routes are reserved exclusively for the use of customers or applications that have contracted for high levels of service.

SUMMARY OF THE INVENTION

A method, system, apparatus, and computer program product is presented for management of a distributed data processing system. A network management framework provides the ability to restrict the use of endpoint resources by receiving a request from a source endpoint for a particular type of action at a target endpoint. A topology mapping is used to derive an endpoint-to-endpoint route for completing the requested action. Other activities along the route are restricted or terminated on an endpoint-by-endpoint basis as necessary in accordance with the derived endpoint-to-endpoint route.

15

BRIEF DESCRIPTION OF THE DRAWINGS

5 The novel features believed characteristic of the invention are set forth in the appended claims. The invention itself, further objectives, and advantages thereof, will be best understood by reference to the following detailed description when read in conjunction
10 with the accompanying drawings, wherein:

Figure 1 is a diagram depicting a known logical configuration of software and hardware resources;

Figure 2A is simplified diagram illustrating a large distributed computing enterprise environment in which the
15 present invention is implemented;

Figure 2B is a block diagram of a preferred system management framework illustrating how the framework functionality is distributed across the gateway and its endpoints within a managed region;

20 **Figure 2C** is a block diagram of the elements that comprise the low cost framework (LCF) client component of the system management framework;

Figure 2D is a diagram depicting a logical configuration of software objects residing within a
25 hardware network similar to that shown in **Figure 2A**;

Figure 2E is a diagram depicting the logical relationships between components within a system management framework that includes two endpoints and a gateway;

Figure 2F is a diagram depicting the logical relationships between components within a system management framework that includes a gateway supporting two DKS-enabled applications;

5 **Figure 2G** is a diagram depicting the logical relationships between components within a system management framework that includes two gateways supporting two endpoints;

10 **Figure 3** is a block diagram depicting components within the system management framework that provide resource leasing management functionality within a distributed computing environment such as that shown in
Figures 2D-2E;

15 **Figure 4** is a block diagram showing data stored by a the IPOP (IP Object Persistence) service;

Figure 5A is a block diagram showing the IPOP service in more detail;

Figure 5B is a network diagram depicting a set of routers that undergo a scoping process;

20 **Figure 6A** is a flowchart depicting a process for obtaining and using an application action object (AAO) within the network management system of the present invention;

25 **Figure 6B** is a flowchart depicting a process for generating an AAO with consideration of whether the requested AAO is directed to a restricted AAO, i.e. an AAO that requires restricted access to endpoints along a route;

30 **Figure 6C** is a flowchart depicting a process for associating an indication of restricted access for endpoints along a route;

Figure 6D is a flowchart depicting a process within a gateway for restricting further usage of an endpoint to which access restrictions are being applied; and

Figure 6E is a flowchart depicting a process for
5 releasing a previously restricted route of endpoints.

DETAILED DESCRIPTION OF THE INVENTION

5 The present invention provides a methodology for managing a distributed data processing system. The manner in which the system management is performed is described further below in more detail after the description of the preferred embodiment of the
10 distributed computing environment in which the present invention operates.

With reference now to **Figure 1**, a diagram depicts a known logical configuration of software and hardware resources. In this example, the software is organized in
15 an object-oriented system. Application object 102, device driver object 104, and operating system object 106 communicate across network 108 with other objects and with hardware resources 110-114.

In general, the objects require some type of
20 processing, input/output, or storage capability from the hardware resources. The objects may execute on the same device to which the hardware resource is connected, or the objects may be physically dispersed throughout a distributed computing environment. The objects request
25 access to the hardware resource in a variety of manners, e.g. operating system calls to device drivers. Hardware resources are generally available on a first-come, first-serve basis in conjunction with some type of arbitration scheme to ensure that the requests for
30 resources are fairly handled. In some cases, priority may be given to certain requesters, but in most

implementations, all requests are eventually processed.

With reference now to **Figure 2A**, the present invention is preferably implemented in a large distributed computer environment **210** comprising up to 5 thousands of "nodes". The nodes will typically be geographically dispersed and the overall environment is "managed" in a distributed manner. Preferably, the managed environment is logically broken down into a series of loosely connected managed regions (MRs) **212**, 10 each with its own management server **214** for managing local resources with the managed region. The network typically will include other servers (not shown) for carrying out other distributed network functions. These include name servers, security servers, file servers, 15 thread servers, time servers and the like. Multiple servers **214** coordinate activities across the enterprise and permit remote management and operation. Each server **214** serves a number of gateway machines **216**, each of which in turn support a plurality of endpoints/terminal 20 nodes **218**. The server **214** coordinates all activity within the managed region using a terminal node manager at server **214**.

With reference now to **Figure 2B**, each gateway machine **216** runs a server component **222** of a system management framework. The server component **222** is a 25 multi-threaded runtime process that comprises several components: an object request broker (ORB) **221**, an authorization service **223**, object location service **225** and basic object adapter (BOA) **227**. Server component **222** 30 also includes an object library **229**. Preferably, ORB **221** runs continuously, separate from the operating system,

and it communicates with both server and client processes through separate stubs and skeletons via an interprocess communication (IPC) facility 219. In particular, a secure remote procedure call (RPC) is used to invoke 5 operations on remote objects. Gateway machine 216 also includes operating system 215 and thread mechanism 217.

The system management framework, also termed distributed kernel services (DKS), includes a client component 224 supported on each of the endpoint machines 10 218. The client component 224 is a low cost, low maintenance application suite that is preferably "dataless" in the sense that system management data is not cached or stored there in a persistent manner.

Implementation of the management framework in this 15 "client-server" manner has significant advantages over the prior art, and it facilitates the connectivity of personal computers into the managed environment. It should be noted, however, that an endpoint may also have an ORB for remote object-oriented operations within the 20 distributed environment, as explained in more detail further below.

Using an object-oriented approach, the system management framework facilitates execution of system management tasks required to manage the resources in the 25 managed region. Such tasks are quite varied and include, without limitation, file and data distribution, network usage monitoring, user management, printer or other resource configuration management, and the like. In a preferred implementation, the object-oriented framework 30 includes a Java runtime environment for well-known advantages, such as platform independence and

standardized interfaces. Both gateways and endpoints operate portions of the system management tasks through cooperation between the client and server portions of the distributed kernel services.

5 In a large enterprise, such as the system that is illustrated in **Figure 2A**, there is preferably one server per managed region with some number of gateways. For a workgroup-size installation, e.g., a local area network, a single server-class machine may be used as both a
10 server and a gateway. References herein to a distinct server and one or more gateway(s) should thus not be taken by way of limitation as these elements may be combined into a single platform. For intermediate size installations, the managed region grows breadth-wise,
15 with additional gateways then being used to balance the load of the endpoints.

The server is the top-level authority over all gateway and endpoints. The server maintains an endpoint list, which keeps track of every endpoint in a managed
20 region. This list preferably contains all information necessary to uniquely identify and manage endpoints including, without limitation, such information as name, location, and machine type. The server also maintains the mapping between endpoints and gateways, and this
25 mapping is preferably dynamic.

As noted above, there are one or more gateways per managed region. Preferably, a gateway is a fully managed node that has been configured to operate as a gateway. In certain circumstances, though, a gateway may be
30 regarded as an endpoint. A gateway always has a network interface card (NIC), so a gateway is also always an endpoint. A gateway usually uses itself as the first

seed during a discovery process. Initially, a gateway does not have any information about endpoints. As endpoints login, the gateway builds an endpoint list for its endpoints. The gateway's duties preferably include:

5 listening for endpoint login requests, listening for endpoint update requests, and (its main task) acting as a gateway for method invocations on endpoints.

As also discussed above, the endpoint is a machine running the system management framework client component,

10 which is referred to herein as a management agent. The management agent has two main parts as illustrated in

Figure 2C: daemon 226 and application runtime library 228. Daemon **226** is responsible for endpoint login and for spawning application endpoint executables. Once an

15 executable is spawned, daemon **226** has no further interaction with it. Each executable is linked with application runtime library **228**, which handles all further communication with the gateway.

Preferably, the server and each of the gateways is a distinct computer. For example, each computer may be a RISC System/6000™ (a reduced instruction set or so-called RISC-based workstation) running the AIX (Advanced

Interactive Executive) operating system. Of course, other machines and/or operating systems may be used as

25 well for the gateway and server machines.

Each endpoint is also a computing device. In one preferred embodiment of the invention, most of the endpoints are personal computers, e.g., desktop machines or laptops. In this architecture, the endpoints need not be high powered or complex machines or workstations. An endpoint computer preferably includes a Web browser such

as Netscape Navigator or Microsoft Internet Explorer. An endpoint computer thus may be connected to a gateway via the Internet, an intranet or some other computer network.

Preferably, the client-class framework running on 5 each endpoint is a low-maintenance, low-cost framework that is ready to do management tasks but consumes few machine resources because it is normally in an idle state. Each endpoint may be "dataless" in the sense that system management data is not stored therein before or 10 after a particular system management task is implemented or carried out.

With reference now to **Figure 2D**, a diagram depicts a logical configuration of software objects residing within a hardware network similar to that shown in **Figure 2A**.

15 The endpoints in **Figure 2D** are similar to the endpoints shown in **Figure 2B**. Object-oriented software, similar to the collection of objects shown in **Figure 1**, executes on the endpoints. Endpoints 230 and 231 support application action object 232 and application object 233, device 20 driver objects 234-235, and operating system objects 236-237 that communicate across a network with other objects and hardware resources.

Resources can be grouped together by an enterprise into managed regions representing meaningful groups.

25 Overlaid on these regions are domains that divide resources into groups of resources that are managed by gateways. The gateway machines provide access to the resources and also perform routine operations on the resources, such as polling. **Figure 2D** shows that 30 endpoints and objects can be grouped into managed regions that represent branch offices 238 and 239 of an

enterprise, and certain resources are controlled by in central office 240. Neither a branch office nor a central office is necessarily restricted to a single physical location, but each represents some of the 5 hardware resources of the distributed application framework, such as routers, system management servers, endpoints, gateways, and critical applications, such as corporate management Web servers. Different types of gateways can allow access to different types of 10 resources, although a single gateway can serve as a portal to resources of different types.

With reference now to **Figure 2E**, a diagram depicts the logical relationships between components within a system management framework that includes two endpoints 15 and a gateway. **Figure 2E** shows more detail of the relationship between components at an endpoint. Network 250 includes gateway 251 and endpoints 252 and 253, which contain similar components, as indicated by the similar reference numerals used in the figure. An endpoint may 20 support a set of applications 254 that use services provided by the distributed kernel services 255, which may rely upon a set of platform-specific operating system resources 256. Operating system resources may include TCP/IP-type resources, SNMP-type resources, and other 25 types of resources. For example, a subset of TCP/IP-type resources may be a line printer (LPR) resource that allows an endpoint to receive print jobs from other endpoints. Applications 254 may also provide self-defined sets of resources that are accessible to 30 other endpoints. Network device drivers 257 send and receive data through NIC hardware 258 to support

communication at the endpoint.

With reference now to **Figure 2F**, a diagram depicts the logical relationships between components within a system management framework that includes a gateway supporting two DKS-enabled applications. Gateway 260 communicates with network 262 through NIC 264. Gateway 260 contains ORB 266 that supports DKS-enabled applications 268 and 269. **Figure 2F** shows that a gateway can also support applications. In other words, a gateway should not be viewed as merely being a management platform but may also execute other types of applications.

With reference now to **Figure 2G**, a diagram depicts the logical relationships between components within a system management framework that includes two gateways supporting two endpoints. Gateway 270 communicates with network 272 through NIC 274. Gateway 270 contains ORB 276 that may provide a variety of services, as is explained in more detail further below. In this particular example, **Figure 2G** shows that a gateway does not necessarily connect with individual endpoints.

Gateway 270 communicates through NIC 278 and network 279 with gateway 280 and its NIC 282. Gateway 280 contains ORB 284 for supporting a set of services. Gateway 280 communicates through NIC 286 and network 287 to endpoint 290 through its NIC 292 and to endpoint 294 through its NIC 296. Endpoint 290 contains ORB 298 while endpoint 294 does not contain an ORB. In this particular example, **Figure 2G** also shows that an endpoint does not necessarily contain an ORB. Hence, any use of endpoint 294 as a resource is performed solely through management

processes at gateway 280.

Figures 2F and 2G also depict the importance of gateways in determining routes/data paths within a highly distributed system for addressing resources within the system and for performing the actual routing of requests for resources. The importance of representing NICs as objects for an object-oriented routing system is described in more detail further below.

As noted previously, the present invention is directed to a methodology for managing a distributed computing environment. A resource is a portion of a computer system's physical units, a portion of a computer system's logical units, or a portion of the computer system's functionality that is identifiable or addressable in some manner to other physical or logical units within the system.

With reference now to **Figure 3**, a block diagram depicts components within the system management framework within a distributed computing environment such as that shown in **Figures 2D-2E**. A network contains gateway 300 and endpoints 301 and 302. Gateway 302 runs ORB 304. In general, an ORB can support different services that are configured and run in conjunction with an ORB. In this case, distributed kernel services (DKS) include Network 20 Endpoint Location Service (NEL) 306, IP Object Persistence (IPOP) service 308, and Gateway Service 310.

The Gateway Service processes action objects, which are explained in more detail below, and directly communicates with endpoints or agents to perform 30 management operations. The gateway receives events from resources and passes the events to interested parties

within the distributed system. The NEL service works in combination with action objects and determines which gateway to use to reach a particular resource. A gateway is determined by using the discovery service of the 5 appropriate topology driver, and the gateway location may change due to load balancing or failure of primary gateways.

Other resource level services may include an SNMP (Simple Network Management Protocol) service that 10 provides protocol stacks, polling service, and trap receiver and filtering functions. The SNMP Service can be used directly by certain components and applications when higher performance is required or the location independence provided by the gateways and action objects 15 is not desired. A Metadata Service can also be provided to distribute information concerning the structure of SNMP agents.

The representation of resources within DKS allows for the dynamic management and use of those resources by 20 applications. DKS does not impose any particular representation, but it does provide an object-oriented structure for applications to model resources. The use of object technology allows models to present a unified appearance to management applications and hide the 25 differences among the underlying physical or logical resources. Logical and physical resources can be modeled as separate objects and related to each other using relationship attributes.

By using objects, for example, a system may 30 implement an abstract concept of a router and then use this abstraction within a range of different router hardware. The common portions can be placed into an

abstract router class while modeling the important differences in subclasses, including representing a complex system with multiple objects. With an abstracted and encapsulated function, the management applications do 5 not have to handle many details for each managed resource. A router usually has many critical parts, including a routing subsystem, memory buffers, control components, interfaces, and multiple layers of communication protocols. Using multiple objects has the 10 burden of creating multiple object identifiers (OIDs) because each object instance has its own OID. However, a first order object can represent the entire resource and contain references to all of the constituent parts.

Each endpoint may support an object request broker, 15 such as ORBs 320 and 322, for assisting in remote object-oriented operations within the DKS environment. Endpoint 301 contains DKS-enabled application 324 that utilizes object-oriented resources found within the distributed computing environment. Endpoint 302 contains 20 target resource provider object or application 326 that services the requests from DKS-enabled application 324. A set of DKS services 330 and 334 support each particular endpoint.

Applications require some type of insulation from 25 the specifics of the operations of gateways. In the DKS environment, applications create action objects that encapsulate command which are sent to gateways, and the applications wait for the return of the action object. Action objects contain all of the information necessary 30 to run a command on a resource. The application does not need to know the specific protocol that is used to

communicate with the resource. The application is unaware of the location of the resource because it issues an action object into the system, and the action object itself locates and moves to the correct gateway. The 5 location independence allows the NEL service to balance the load between gateways independently of the applications and also allows the gateways to handle resources or endpoints that move or need to be serviced by another gateway.

10 The communication between a gateway and an action object is asynchronous, and the action objects provide error handling and recovery. If one gateway goes down or becomes overloaded, another gateway is located for executing the action object, and communication is 15 established again with the application from the new gateway. Once the controlling gateway of the selected endpoint has been identified, the action object will transport itself there for further processing of the command or data contained in the action object. If it is 20 within the same ORB, it is a direct transport. If it is within another ORB, then the transport can be accomplished with a "Moveto" command or as a parameter on a method call.

Queuing the action object on the gateway results in 25 a controlled process for the sending and receiving of data from the IP devices. As a general rule, the queued action objects are executed in the order that they arrive at the gateway. The action object may create child action objects if the collection of endpoints contains 30 more than a single ORB ID or gateway ID. The parent action object is responsible for coordinating the completion status of any of its children. The creation

of child action objects is transparent to the calling application. A gateway processes incoming action objects, assigns a priority, and performs additional security challenges to prevent rogue action object attacks. The action object is delivered to the gateway that must convert the information in the action object to a form suitable for the agent. The gateway manages multiple concurrent action objects targeted at one or more agents, returning the results of the operation to the calling managed object as appropriate.

In the preferred embodiment, potentially leasable target resources are Internet protocol (IP) commands, e.g. pings, and Simple Network Management Protocol (SNMP) commands that can be executed against endpoints in a managed region. Referring again to **Figures 2F** and **2G**, each NIC at a gateway or an endpoint may be used to address an action object. Each NIC is represented as an object within the IPOP database, which is described in more detail further below.

The Action Object IP (AOIP) Class is a subclass of the Action Object Class. AOIP objects are the primary vehicle that establishes a connection between an application and a designated IP endpoint using a gateway or stand-alone service. In addition, the Action Object SNMP (AOSnmp) Class is also a subclass of the Action Object Class. AOSnmp objects are the primary vehicle that establishes a connection between an application and a designated SNMP endpoint via a gateway or the Gateway Service. However, the present invention is primarily concerned with IP endpoints.

The AOIP class should include the following: a constructor to initialize itself; an interface to the NEL

service; a mechanism by which the action object can use the ORB to transport itself to the selected gateway; a mechanism by which to communicate with the SNMP stack in a stand-alone mode; a security check verification of

5 access rights to endpoints; a container for either data or commands to be executed at the gateway; a mechanism by which to pass commands or classes to the appropriate gateway or endpoint for completion; and public methods to facilitate the communication between objects.

10 The instantiation of an AOIP object creates a logical circuit between an application and the targeted gateway or endpoint. This circuit is persistent until command completion through normal operation or until an exception is thrown. When created, the AOIP object

15 instantiates itself as an object and initializes any internal variables required. An action object IP may be capable of running a command from inception or waiting for a future command. A program that creates an AOIP object must supply the following elements: address of

20 endpoints; function to be performed on the endpoint, class, or object; and data arguments specific to the command to be run. A small part of the action object must contain the return end path for the object. This may identify how to communicate with the action object in

25 case of a breakdown in normal network communications. An action object can contain either a class or object containing program information or data to be delivered eventually to an endpoint or a set of commands to be performed at the appropriate gateway. Action objects IP 30 return back a result for each address endpoint targeted.

Using commands such as "Ping", "Trace Route", "Wake-On LAN", and "Discovery", the AOIP object performs

DRAFT EDITION 1

the following services: facilitates the accumulation of metrics for the user connections; assists in the description of the topology of a connection; performs Wake-On LAN tasks using helper functions; and discovers 5 active agents in the network environment.

The NEL service finds a route (data path) to communicate between the application and the appropriate endpoint. The NEL service converts input to protocol, network address, and gateway location for use by action 10 objects. The NEL service is a thin service that supplies information discovered by the IPOP service. The primary roles of the NEL service are as follows: support the requests of applications for routes; maintain the gateway and endpoint caches that keep the route information; 15 ensure the security of the requests; and perform the requests as efficiently as possible to enhance performance.

For example, an application requires a target endpoint (target resource) to be located. The target is 20 ultimately known within the DKS space using traditional network values, i.e. a specific network address and a specific protocol identifier. An action object is generated on behalf of an application to resolve the network location of an endpoint. The action object asks 25 the NEL service to resolve the network address and define the route to the endpoint in that network.

One of the following is passed to the action object to specify a destination endpoint: an EndpointAddress object; a fully decoded NetworkAddress object; and a 30 string representing the IP address of the IP endpoint. In combination with the action objects, the NEL service determines which gateway to use to reach a particular

resource. The appropriate gateway is determined using the discovery service of the appropriate topology driver and may change due to load balancing or failure of primary gateways. An "EndpointAddress" object must

5 consist of a collection of at least one or more unique managed resource IDs. A managed resource ID decouples the protocol selection process from the application and allows the NEL service to have the flexibility to decide the best protocol to reach an endpoint. On return from

10 the NEL service, an "AddressEndpoint" object is returned, which contains enough information to target the best place to communicate with the selected IP endpoints. It should be noted that the address may include

15 protocol-dependent addresses as well as protocol-independent addresses, such as the virtual private network id and the IPOP Object ID. These additional addresses handle the case where duplicate addresses exist in the managed region.

When an action needs to be taken on a set of

20 endpoints, the NEL service determines which endpoints are managed by which gateways. When the appropriate gateway is identified, a single copy of the action object is distributed to each identified gateway. The results from the endpoints are asynchronously merged back to the

25 caller application through the appropriate gateways. Performing the actions asynchronously allows for tracking all results whether the endpoints are connected or disconnected. If the action object IP fails to execute an action object on the target gateway, NEL is consulted

30 to identify an alternative path for the command. If an alternate path is found, the action object IP is transported to that gateway and executed. It may be

assumed that the entire set of commands within one action object IP must fail before this recovery procedure is invoked.

With reference now to **Figure 4**, a block diagram 5 shows the manner in which data is stored by the IPOP (IP Object Persistence) service. IPOP service database 402 contains endpoint database table 404, system database table 406, and network database table 408. Each table 10 contains a set of topological (topo) objects for facilitating the reservation of resources at IP endpoints and the execution of action objects. Information within IPOP service database 402 allows applications to generate action objects for resources previously identified as IP objects through a discovery process across the 15 distributed computing environment. **Figure 4** merely shows that the topo objects may be separated into a variety of categories that facilitate processing on the various objects. The separation of physical network categories facilitates the efficient querying and storage of these 20 objects while maintaining the physical network relationships in order to produce a graphical user interface of the network topology.

With reference now to **Figure 5A**, a block diagram shows the IPOP service in more detail. In the preferred 25 embodiment of the present invention, an IP driver subsystem is implemented as a collection of software components for discovering, i.e. detecting, IP "objects", i.e. IP networks, IP systems, and IP endpoints by using physical network connections. This discovered physical 30 network is used to create topology data that is then provided through other services via topology maps

accessible through a graphical user interface (GUI) or for the manipulation of other applications. The IP driver system can also monitor objects for changes in IP topology and update databases with the new topology 5 information. The IPOP service provides services for other applications to access the IP object database.

IP driver subsystem 500 contains a conglomeration of components, including one or more IP drivers 502. Every IP driver manages its own "scope", which is described in 10 more detail further below, and every IP driver is assigned to a topology manager within Topology Service 504, which can serve may than one IP driver. Topology Service 504 stores topology information obtained from discovery controller 506. The information stored within 15 the Topology Service may include graphs, arcs, and the relationships between nodes determined by IP mapper 508. Users can be provided with a GUI to navigate the topology, which can be stored within a database within the Topology Service.

20 IPOP service 510 provides a persistent repository 511 for discovered IP objects. Persistent repository 511 also stores application action object restricted session identifiers 512 and physical network topology route 513, as discussed in more detail below with respect to **Figures** 25 **6A-6E**. Discovery controller 506 detects IP objects in Physical IP networks 514, and monitor controller 516 monitors IP objects. A persistent repository, such as IPOP database 511, is updated to contain information about the discovered and monitored IP objects. IP driver 30 may use temporary IP data store component 518 and IP data cache component 520 as necessary for caching IP objects

or storing IP objects in persistent repository 511, respectively. As discovery controller 506 and monitor controller 516 perform detection and monitoring functions, events can be written to network event manager 5 application 522 to alert network administrators of certain occurrences within the network, such as the discovery of duplicate IP addresses or invalid network masks.

External applications/users 524 can be other users, 10 such as network administrators at management consoles, or applications that use IP driver GUI interface 526 to configure IP driver 502, manage/unmanage IP objects, and manipulate objects in persistent repository 512. Configuration service 528 provides configuration 15 information to IP driver 502. IP driver controller 532 serves as central control of all other IP driver components.

Referring back to **Figure 2G**, a network discovery engine is a distributed collection of IP drivers that are 20 used to ensure that operations on IP objects by gateways 260, 270, and 280 can scale to a large installation and provide fault-tolerant operation with dynamic start/stop or reconfiguration of each IP driver. The IPOP Service manages discovered IP objects; to do so, the IPOP Service 25 uses a distributed database in order to efficiently service query requests by a gateway to determine routing, identity, or a variety of details about an endpoint. The IPOP Service also services queries by the Topology Service in order to display a physical network or map 30 them to a logical network, which is a subset of a physical network that is defined programmatically or by

an administrator. IPOP fault tolerance is also achieved by distribution of IPOP data and the IPOP Service among many Endpoint ORBs.

The IP Monitor Controller uses SNMP polls to 5 determine if there have been any configuration changes in an IP system. It also looks for any IP endpoints added to or deleted from an IP system. The IP Monitor Controller also monitors the statuses of IP endpoints in an IP system. In order to reduce network traffic, an IP 10 driver will use SNMP to get the status of all IP endpoints in an IP system in one query unless an SNMP agent is not running on the IP system. Otherwise, an IP driver will use "Ping" instead of SNMP. An IP driver will use "Ping" to get the status of an IP endpoint if it 15 is the only IP endpoint in the system since the response from "Ping" is quicker than SNMP.

One or more IP drivers can be deployed to provide distribution of IP discovery and promote scalability of IP driver subsystem services in large networks where a 20 single IP driver subsystem is not sufficient to discover and monitor all IP objects. Each IP discovery driver performs discovery and monitoring on a collection of IP resources within the driver's "scope". A driver's scope, which is explained in more detail below, is simply the 25 set of IP subnets for which the driver is responsible for discovering and monitoring. Network administrators generally partition their networks into as many scopes as needed to provide distributed discovery and satisfactory performance.

30 A potential risk exists if the scope of one driver overlaps the scope of another, i.e., if two drivers attempt to discover/monitor the same device. Accurately

defining unique and independent scopes may require the development of a scope configuration tool to verify the uniqueness of scope definitions. Routers also pose a potential problem in that while the networks serviced by 5 the routers will be in different scopes, a convention needs to be established to specify to which network the router "belongs", thereby limiting the router itself to the scope of a single driver.

Some ISPs may have to manage private networks whose 10 addresses may not be unique across the installation, like 10.0.0.0 network. In order to manage private networks properly, first, the IP driver has to be installed inside the internal networks in order to be able to discover and manage the networks. Second, since the discovered IP 15 addresses may not be unique in across an entire installation that consists of multiple regions, multiple customers, etc., a private network ID has to be assigned to the private network addresses. In the preferred embodiment, the unique name of a subnet becomes 20 "privateNetworkId\subnetAddress". Those customers that do not have duplicate networks address can just ignore the private network ID; the default private network ID is 0.

If Network Address Translator (NAT) is installed to 25 translate the internal IP addresses to Internet IP addresses, users can install the IP drivers outside of NAT and manage the IP addresses inside the NAT. In this case, an IP driver will see only the translated IP addresses and discover only the IP addresses translated. If not all IP addresses inside the NAT are translated, an 30 IP driver will not able to discover all of them. However, if IP drivers are installed this way, users do not have to configure the private network ID.

Scope configuration is important to the proper operation of the IP drivers because IP drivers assume that there are no overlaps in the drivers' scopes. Since there should be no overlaps, every IP driver has complete control over the objects within its scope. A particular IP driver does not need to know anything about the other IP drivers because there is no synchronization of information between IP drivers. The Configuration Service provides the services to allow the DKS components to store and retrieve configuration information for a variety of other services from anywhere in the networks. In particular, the scope configuration will be stored in the Configuration Services so that IP drivers and other applications can access the information.

The ranges of addresses that a driver will discover and monitor are determined by associating a subnet address with a subnet mask and associating the resulting range of addresses with a subnet priority. An IP driver is a collection of such ranges of addresses, and the subnet priority is used to help decide the system address. A system can belong to two or more subnets, such as is commonly seen with a Gateway. The system address is the address of one of the NICs that is used to make SNMP queries. A user interface can be provided, such as an administrator console, to write scope information into the Configuration Service. System administrators do not need to provide this information at all, however, as the IP drivers can use default values.

An IP driver gets its scope configuration information from the Configuration Service, which may be stored using the following format:

DRAFT

```
scopeID=driverID,anchorname,subnetAddress:subnetMask [  
:privateNetworkId:privateNetworkName:subnetPriority] [,  
subnetAddress:subnetMask:privateNetworkId:privateNetworkN  
ame:subnetPriority]]
```

5

Typically, one IP driver manages only one scope. Hence, the "scopeID" and "driverID" would be the same. However, the configuration can provide for more than one scope managed by the same driver. "Anchorname" is the 10 name in the name space in which the Topology Service will put the IP networks objects.

A scope does not have to include an actual subnet configured in the network. Instead, users/administrators can group subnets into a single, logical scope by

15 applying a bigger subnet mask to the network address. For example, if a system has subnet "147.0.0.0" with mask of "255.255.0.0" and subnet "147.1.0.0" with a subnet mask of "255.255.0.0", the subnets can be grouped into a single scope by applying a mask of "255.254.0.0". Assume 20 that the following table is the scope of IP Driver 2.

The scope configuration for IP Driver 2 from the Configuration Service would be:

```
2=2,ip,147.0.0.0:255.254.0.0,146.100.0.0:255.255.0.0,  
69.0.0.0:255.0.0.0.
```

25

Subnet address	Subnet mask
147.0.0.0	255.255.0.0
147.1.0.0	255.255.0.0
146.100.0.0	255.255.0.0
69.0.0.0	255.0.0.0

In general, an IP system is associated with a single IP address, and the "scoping" process is a straightforward association of a driver's ID with the system's IP address.

5 Routers and multi-homed systems, however, complicate the discovery and monitoring process because these devices may contain interfaces that are associated with different subnets. If all subnets of routers and multi-homed systems are in the scope of the same driver,
10 the IP driver will manage the whole system. However, if the subnets of routers and multi-homed systems are across the scopes of different drivers, a convention is needed to determine a dominant interface: the IP driver that manages the dominant interface will manage the router
15 object so that the router is not being detected and monitored by multiple drivers; each interface is still managed by the IP driver determined by its scope; the IP address of the dominant interface will be assigned as the system address of the router or multi-homed system; and
20 the smallest (lowest) IP address of any interface on the router will determine which driver includes the router object within its scope.

Users can customize the configuration by using the subnet priority in the scope configuration. The subnet
25 priority will be used to determinate the dominant interface before using the lowest IP address. If the subnet priorities are the same, the lowest IP address is then used. Since the default subnet priority would be "0", then the lowest IP address would be used by default.

30 With reference now to **Figure 5B**, a network diagram depicts a network with a router that undergoes a scoping process. IP driver D1 will include the router in its

scope because the subnet associated with that router interface is lower than the other three subnet addresses. However, each driver will still manage those interfaces inside the router in its scope. Drivers D2 and D3 will 5 monitor the devices within their respective subnets, but only driver D1 will store information about the router itself in the IPOP database and the Topology Service database.

If driver D1's entire subnet is removed from the 10 router, driver D2 will become the new "owner" of the router object because the subnet address associated with driver D2 is now the lowest address on the router. Because there is no synchronization of information between the drivers, the drivers will self-correct over 15 time as they periodically rediscover their resources. When the old driver discovers that it no longer owns the router, it deletes the router's information from the databases. When the new driver discovers the router's lowest subnet address is now within its scope, the new 20 driver takes ownership of the router and updates the various data bases with the router's information. If the new driver discovers the change before the old driver has deleted the object, then the router object may be briefly represented twice until the old owner deletes the 25 original representation.

There are two kinds of associations between IP objects. One is "IP endpoint in IP system" and the other is "IP endpoint in IP network". The implementation of associations relies on the fact that an IP endpoint has 30 the object IDs (OIDs) of the IP system and the IP network in which it is located. Based on the scopes, an IP driver can partition all IP networks, IP Systems, and IP

endpoints into different scopes. A network and all its IP endpoints will always be assigned in the same scope.

However, a router may be assigned to an IP Driver, but some of its interfaces are assigned to different to

5 different IP drivers. The IP drivers that do not manage the router but manage some of its interfaces will have to create interfaces but not the router object. Since those IP drivers do not have a router object ID to assign to its managed interfaces, they will assign a unique system
10 name instead of object ID in the IP endpoint object to provide a link to the system object in a different driver. Because of the inter-scope association, when the IP Persistence Service (IPOP) is queried to find all the IP endpoints in system, it will have to search not only
15 IP endpoints with the system ID but also IP endpoints with its system name. If a distributed IP Persistence Service is implemented, the IP Persistence Service has to provide extra information for searching among IP Persistence Services.

20 As noted above, in the DKS environment, an application requests the creation of an action object that encapsulates a command that is sent to a gateway, and the application waits for the return of the action object's completion. Action objects generally contain
25 all of the information necessary to run a command on a resource. The application does not necessarily need to know the specific protocol that is used to communicate with the resource. Moreover, the application may be unaware of the location of the resource because it issues
30 an action object into the system, and the action object itself locates and moves to the correct gateway.

For example, an application requires a target

resource (target endpoint) to be located. The target object is ultimately known within the DKS space using traditional network values, i.e. a specific network address and a specific protocol identifier. However, an application can address a target object with an Object ID. An action object is generated on behalf of an application to resolve the network location of an endpoint. The action object asks the NEL service to resolve the network address and define the route to the endpoint in that network. One benefit of location independence is that the NEL service can balance the load between gateways independently of the applications and also allows the gateways to handle resources or endpoints that move or need to be serviced by another gateway.

In order to fulfill quality-of-service guarantees within a network management system, which might consist of a million devices or more, a service provider may require performance measurements along various network routes throughout the system. A performance measurement operation on a target resource may be routed through a particular node in the network, and the node may be used simultaneously by multiple activities, including application-specific activities and other management activities. Any of these activities may consume a significant amount of the node's resources, such as its bandwidth, thereby affecting measured response times. Hence, those other activities should be restricted or terminated during a performance measurement operation. To do so, computational resources throughout the system need to be controllable so that the management system can obtain accurate performance measurements along particular routes that are independent of other activities that are

occurring along the route.

In order to address the issues of exclusive access to resources, the present invention restricts consumption of resources along particular network routes so that a 5 network management system can accurately obtain performance measurements throughout the highly distributed system. Moreover, the network management system can restrict the usage of routes for other purposes. For example, a service provider can contract 10 with customers to provide a high level of service that requires exclusive use of a particular route. These features are explained in more detail below with respect to **Figures 6A-6E**.

With respect to **Figure 6A**, a flowchart depicts a 15 process for obtaining and using an application action object (AAO) within the network management system of the present invention. An application action object is a class of objects that extends an action object class in a manner that is appropriate for a particular application. 20 The process begins when an application requests an application action object for a target endpoint from the Gateway service (step 602).

It should be noted that the process shown in **Figure 6A** is generic with respect to an application requesting 25 and obtaining action objects. However, given the processing context shown in **Figures 6B-6E**, it may be assumed that the requested AAO in step 602 is a special type of AAO that requires that the network management system execute the AAO with a high level of performance, 30 in which case the network management system applies access restrictions to endpoints along the route that

will be used by the AAO.

The present invention provides a methodology through which a network management framework can restrict access to endpoints along logical routes through a network. The 5 initiation of restricted access begins with a request for a special type of AAO that requires restricted access for its proper implementation or proper execution. Depending upon the system implementation, there may be several types of action objects for which DKS will automatically 10 initiate the reservation of a restricted route.

For example, one special type of application action object would be an action that executes a performance measurement; in general, an accurate performance measurement requires that endpoints along a network route 15 not be used by other action objects while the action object for the performance measurement is being processed or executed. Hence, if a particular request for an instance of an action object is defined as a type of action object related to performance measurements, then 20 the processes shown in **Figure 6B-6E** may be engaged.

Referring again to **Figure 6A**, the process continues when the Gateway Service asks the NEL service to decode the target endpoint from the request (step 604). As noted previously, one of the primary roles of the NEL 25 service is to support the requests from applications for routes, as explained above with respect to **Figure 3**. The NEL service then asks the IPOP service to decode the endpoint object (step 606). Assuming that the processing has been successfully accomplished, IPOP returns an 30 appropriate AAO to the NEL service (step 608), and the NEL service returns the AAO to the Gateway service (step

610). The Gateway service then returns the AAO to the application (step 612). The application then performs the desired action (step 614), such as a performance measurement for an endpoint-to-endpoint route, and the 5 process is complete. As is apparent with respect to **Figure 6A**, an application action object that may require special route restrictions can be processed by an application in a manner similar to that used to process any other type of application action object.

10 With respect to **Figure 6B**, a flowchart depicts a process for generating an AAO with consideration of whether the requested AAO is directed to a restricted AAO, i.e. an AAO that requires restricted access to endpoints along a route. **Figure 6B** provides more detail 15 for step 608 shown in **Figure 6A**. The process begins with IPOP receiving the target endpoint from the NEL service and determining whether another application has already requested that access to or use of the target endpoint should be restricted, which may be indicated by setting a 20 restrict flag to "true" (step 620). If so, then the request for the AAO is rejected (step 622), and the process is complete. In this case, the target endpoint has an associated indication that it is already being used exclusively by another application.

25 If the restrict flag is not set for the target endpoint, then a determination is made as to whether the requested AAO is a type of AAO that requires a restricted route for its proper completion or execution (step 624). In other words, if the requested AAO were to be granted, 30 then the AAO should be processed or executed at a high level of service, performance, or priority. If not, then

IPOP has determined that it is clear to process the request. IPOP decodes the AAO address from the IPOPOid and populates the AAO with the required information (step 626). IPOP then returns the AAO to the NEL service (step 5 628), and the process is complete.

If a restricted route is required for the proper completion or execution of the action object, then IPOP restricts the route for the requested AAO (step 630), and completes the creation of the AAO at steps 626-628, and 10 the process is complete. As noted above, IPOP may examine the type of requested AAO in order to determine whether or not execution of the AAO would require a high level of performance such that a restricted route should be reserved for the application action object's 15 execution. As noted above with respect to **Figure 3**, the instantiation of an action object creates a logical circuit between a source endpoint and a target endpoint; various endpoints along a route through a network are used by the network management framework to complete the 20 execution of an action object. The present invention may be viewed as a methodology for reserving the logical circuit; the logical circuit is reserved if the type of requested action object requires a reserved logical circuit or if, for other reasons, the network management 25 system should complete or should attempt to complete the requested action object with a high level of service or performance. The network management system may have other reasons for using a reserved route of endpoints, such as a customer guarantee in which a service provider 30 has contracted to provide a particular quality of service.

It should also be noted that IPOP may determine within the processing of step 630 that a particular endpoint within the computed route for the requested AAO may already have been granted restricted access. Hence, 5 IPOP would be required to select a new route for the requested AAO that does not contain the restricted endpoint before IPOP could restrict an entire route. Otherwise, the request for the AAO may still be denied if an entire route for the requested AAO cannot be secured 10 because all possible routes for the requested AAO contain at least one endpoint that has already been reserved with restricted access.

With respect to **Figure 6C**, a flowchart depicts a process for associating an indication of restricted 15 access for endpoints along a route. **Figure 6C** provides more detail for step 630 shown in **Figure 6B**. The process begins with IPOP determining a logical route through the distributed system along a series of endpoints from the application endpoint or source endpoint to the target 20 endpoint as required to execute the requested AAO (step 642). IPOP then defines a restricted session number that will be associated with this route and stores it in the AAO (step 644), and IPOP gets the set of all endpoints for the route (step 646).

25 A next endpoint from the set of endpoints in the route is obtained (step 648), and IPOP sets the restrict flag for the endpoint that is currently being processed (step 650). IPOP also notifies the Gateway service to reject further usage of the current endpoint by other 30 applications (step 652). A determination is then made whether there are other endpoints in the set of endpoints

for the route (step 654), and if so, then the process branches back to step 648 to process another endpoint. Otherwise, the process of restricting the route is then complete.

5 With respect to **Figure 6D**, a flowchart depicts a process within a gateway for restricting further usage of an endpoint to which access restrictions are being applied. **Figure 6D** provides more detail for step 652 shown in **Figure 6C**. A gateway receives the notification
10 to restrict use of an endpoint, and the notification includes a restricted session number (step 660). The endpoint in the notification is within the route for the requested AAO, as determined by IPOP, as explained with respect to **Figure 6C**, and the gateway that processes the
15 notification is responsible for managing the endpoint. The gateway then terminates all other application action objects that have been using the endpoint for which restricted access has been granted to the other application action object (step 662), and the process is
20 complete.

With respect to **Figure 6E**, a flowchart depicts a process for releasing a previously restricted route of endpoints. **Figure 6E** provides more detail for steps that would occur after an application had used a route that
25 had been restricted using processes shown in **Figures 6A-6D**. The process begins with the application notifying the gateway service that activities have been completed for this particular AAO with its associated restricted session number (step 672). The gateway then informs IPOP
30 that the restricted route is no longer required and passes the AAO to the IPOP service (step 674). IPOP

fetches the restricted session number from the AAO (step 676), and IPOP fetches all endpoints along the route that is associated with the restricted session number (step 678).

5 A next endpoint from the set of endpoints in the route is obtained (step 680), and IPOP resets the restrict flag for the endpoint that is currently being processed (step 682). IPOP also notifies the Gateway service to reject further usage of the current endpoint
10 by other applications, if necessary. A determination is then made whether there are other endpoints in the set of endpoints for the route (step 684), and if so, then the process branches back to step 680 to process another endpoint. Otherwise, the process of releasing the
15 restrictions on the endpoints along the route is then complete.

The advantages of the present invention should be apparent in view of the detailed description of the invention that is provided above. In prior art systems,
20 restricting access to endpoints has generally required manual configuration of all possible applications using a given endpoint. Hence, it has been difficult to obtain accurate quality-of-service measurements at the application level; prior art systems have derived
25 independent measurements from network latency or network response times and operating system level response times.

The present invention provides a network management framework that has the ability to restrict the use of endpoint resources along network routes after receiving,
30 from a source endpoint, certain types of requests for actions at a target endpoint. A topology mapping is used

to derive an endpoint-to-endpoint route for completing the requested action. Other activities along the route are restricted or terminated on an endpoint-by-endpoint basis as necessary in accordance with the derived 5 endpoint-to-endpoint route. Consumption of resources along particular network routes might be restricted that the performance of various target resources can be accurately measured.

In addition, particular routes could be restricted 10 so that they are reserved exclusively for the use of customers or applications that have contracted for high levels of service. A service provider can offer scaled compensation models to customers such that the customers would be charged based on various levels of service, 15 including exclusive use of resources for variable periods of time if necessary. The network management system of the present invention allows QOS guarantees to be provided and verified.

The present invention is not limited to restricting 20 physical routes through a network; physical networks are controlled as logical networks of endpoints. Hence, the present invention could restrict a route between a source endpoint and a target endpoint that are the same endpoint; in effect, the source endpoint could isolate 25 itself in certain aspects when necessary.

It is important to note that while the present invention has been described in the context of a fully functioning data processing system, those of ordinary skill in the art will appreciate that the processes of 30 the present invention are capable of being distributed in the form of instructions in a computer readable medium and a variety of other forms, regardless of the

particular type of signal bearing media actually used to carry out the distribution. Examples of computer readable media include media such as EPROM, ROM, tape, paper, floppy disc, hard disk drive, RAM, and CD-ROMs and 5 transmission-type media, such as digital and analog communications links.

The description of the present invention has been presented for purposes of illustration but is not intended to be exhaustive or limited to the disclosed 10 embodiments. Many modifications and variations will be apparent to those of ordinary skill in the art. The embodiments were chosen to explain the principles of the invention and its practical applications and to enable others of ordinary skill in the art to understand the 15 invention in order to implement various embodiments with various modifications as might be suited to other contemplated uses.